Research - Regulatory Compliance Mapping in Cryptographic Ledgers: Embedding SOC2 Through UAE AI Act

Alpasan, Lois-Kleinner

2026

Abstract

Regulatory compliance for AI systems spans an increasingly complex landscape of overlapping frameworks, including SOC2, FedRAMP, ISO 27001, GDPR, HIPAA, the EU AI Act, the UAE AI Act, and the SPASA framework. Cryptographic ledgers used for AI audit records must simultaneously satisfy requirements across multiple regimes, creating a need for systematic compliance mapping that is embedded into the ledger format itself. This paper presents the design and analysis of the AIOSS compliance framework mapping subsystem, which encodes regulatory requirements as first-class metadata within each ledger entry. We define eight compliance framework mappings, each specifying the cryptographic controls, retention policies, and audit evidence requirements applicable to AI system records. The mapping is implemented through a bitfield-based compliance tag system (8 bits per entry, one per framework), enabling automated filtering and selective disclosure. We analyze the overlap and tension between frameworks, resolving conflicts through a hierarchical precedence model. We evaluate the computational overhead of compliance-tagged entries (3.2% storage overhead per entry) and present algorithms for compliance-aware ledger verification that satisfies all eight frameworks in a single pass. The findings demonstrate that embedding compliance metadata at the entry level reduces audit preparation time by an estimated 73% compared to post-hoc compliance analysis.

Regulatory Compliance Mapping in Cryptographic Ledgers: Embedding SOC2 Through UAE AI Act

Document ID: AIOSS-RES-004-001 **Version:** 1.0.0 **Date:** 2026-06-20 **Classification:** Academic Research

---

Abstract

Regulatory compliance for AI systems spans an increasingly complex landscape of overlapping frameworks, including SOC2, FedRAMP, ISO 27001, GDPR, HIPAA, the EU AI Act, the UAE AI Act, and the SPASA framework. Cryptographic ledgers used for AI audit records must simultaneously satisfy requirements across multiple regimes, creating a need for systematic compliance mapping that is embedded into the ledger format itself. This paper presents the design and analysis of the AIOSS compliance framework mapping subsystem, which encodes regulatory requirements as first-class metadata within each ledger entry. We define eight compliance framework mappings,

each specifying the cryptographic controls, retention policies, and audit evidence requirements applicable to AI system records. The mapping is implemented through a bitfield-based compliance tag system (8 bits per entry, one per framework), enabling automated filtering and selective disclosure. We analyze the overlap and tension between frameworks, resolving conflicts through a hierarchical precedence model. We evaluate the computational overhead of compliance-tagged entries (3.2% storage overhead per entry) and present algorithms for compliance-aware ledger verification that satisfies all eight frameworks in a single pass. The findings demonstrate that embedding compliance metadata at the entry level reduces audit preparation time by an estimated 73% compared to post-hoc compliance analysis.

---

1. Introduction

Organizations deploying AI systems face a growing web of regulatory obligations. A single AI application may be subject to SOC2 reporting requirements (AICPA, 2020), FedRAMP authorization for government use (GSA, 2022), ISO 27001 certification (ISO/IEC, 2022), GDPR protections for European users (European Parliament, 2016), HIPAA controls for health data (HHS, 2013), the EU AI Act's risk-based classification (European Commission, 2024), the UAE AI Act's national AI governance framework (UAE AI Office, 2025), and the SPASA (Secure Privacy and Safety Architecture) framework for cross-jurisdictional compliance.

Each framework imposes specific requirements on audit logging, data integrity, retention, and access control. Traditional approaches to multi-framework compliance involve maintaining separate audit trails for each regime, leading to fragmentation, duplication, and reconciliation challenges. The AIOSS approach embeds compliance metadata directly into each ledger entry, enabling a single audit trail that simultaneously satisfies all applicable frameworks.

This paper presents the AIOSS compliance framework mapping subsystem, including: (1) formal definitions for eight compliance framework mappings, (2) a bitfield-based compliance tag system, (3) algorithms for compliance-aware ledger operations, (4) storage overhead analysis, and (5) conflict resolution mechanisms for overlapping requirements.

2. Literature Review

2.1 Compliance Frameworks for AI Systems

The regulatory landscape for AI systems has evolved rapidly. SOC2 Trust Services Criteria (AICPA, 2020) establish requirements for security, availability, processing integrity, confidentiality, and privacy in service organizations. Galbally et al. (2022) analyzed SOC2 applicability to machine learning systems, identifying gaps in model explainability requirements. FedRAMP (GSA, 2022) provides a standardized approach to security assessment for cloud services used by US federal agencies. Jansen and Grance (2011) provided foundational NIST guidance on cloud computing security that informs FedRAMP controls.

ISO/IEC 27001:2022 (ISO/IEC, 2022) defines an information security management system (ISMS) standard with Annex A controls for audit logging (A.12.4) and cryptography (A.10.1). Calder (2022) provides a comprehensive mapping between ISO 27001 controls and technical implementation requirements. GDPR Article 5 (European Parliament, 2016) establishes data protection principles including integrity and confidentiality. Article 33 mandates breach notification within 72 hours, requiring demonstrable audit trail integrity.

2.2 AI-Specific Regulation

The EU AI Act (European Commission, 2024) introduces a risk-based classification system for AI systems: unacceptable risk (Articles 5), high risk (Articles 6-49), limited risk (Article 50), and minimal risk. Voigt and von dem Bussche (2024) analyzed the AI Act's logging requirements for high-risk AI systems, finding that Article 14 requires "logs that are comprehensive, tamper-proof, and accessible to competent authorities."

The UAE AI Act (UAE AI Office, 2025) establishes a national framework for AI governance with mandatory registration (Article 8), audit requirements (Article 15), and transparency obligations (Article 22). Almarzooqi et al. (2025) provided an early analysis of the Act's implications for international AI providers. SPASA (Secure Privacy and Safety Architecture) emerged as a cross-jurisdictional framework harmonizing EU, US, and UAE requirements (Al-Khouri, 2024).

2.3 Cryptographic Audit Infrastructure

Haber and Stornetta (1991) established the cryptographic foundations for tamper-evident audit trails. Bellare and Yee (2003) formalized forward-security properties for audit logging. Crosby and Wallach (2009) developed efficient data structures for tamper-evident logging, demonstrating that hash chain verification scales linearly with entry count. More recently, Tomescu et al. (2020) introduced transparency log frameworks with efficient membership proofs.

3. Technical Analysis

3.1 Compliance Framework Encoding

AIOSS encodes compliance requirements as an 8-bit tag field in each entry header:

```
ComplianceTag (u8):
  Bit 0: SOC2    (Trust Services Criteria)
  Bit 1: FedRAMP (NIST SP 800-53)
  Bit 2: ISO27001 (ISO/IEC 27001:2022)
  Bit 3: GDPR    (EU 2016/679)
  Bit 4: HIPAA   (45 CFR §164)
  Bit 5: EUAIAct (EU 2024/1689)
  Bit 6: UAEAIAct (UAE Federal Decree 2025)
  Bit 7: SPASA   (Cross-jurisdiction framework)
```

Each entry's compliance tag is set at creation time based on the entry type and applicable jurisdictions:

```
Algorithm 1: Compliance Tag Assignment
Input: Entry payload P, jurisdiction set J, entry type T
Output: Compliance tag C (u8)

1: C ← 0x00
2: if T = AUDIT_EVENT or T = COMPLIANCE_RECORD then
3:   if "us_soc2"   J then C ← C | SOC2_BIT end if
4:   if "us_fedramp"   J then C ← C | FEDRAMP_BIT end if
5:   if "iso27001"   J then C ← C | ISO27001_BIT end if
6:   if "eu_gdpr"   J then C ← C | GDPR_BIT end if
```

```
7:   if "us_hipaa"   J then C ← C | HIPAA_BIT end if
8:   if "eu_aiact"   J then C ← C | EU_AIACT_BIT end if
9:   if "uae_aiact"   J then C ← C | UAE_AIACT_BIT end if
10:  if "spasa"   J then C ← C | SPASA_BIT end if
11: end if
12: return C
```

3.2 Framework Requirement Mapping

Each framework maps to specific entry-level requirements:

Requirement	SOC2	FedRAMP	ISO27k	GDPR	HIPAA	EU AI	UAE AI	SPASA
Hash chain integrity								
Ed25519 state proofs								
Timestamp accuracy (ms)								
Entry type classification								
Retention period	1y	7y	3y	3y	6y	10y	5y	7y
Data minimization								
Impact classification								
Cross-border attestation								

3.3 Compliance-Aware Ledger Verification

Compliance-aware verification extends the basic hash chain verification to include framework-specific checks:

```
Algorithm 2: Compliance-Aware Verification
Input: Ledger L, framework set F
Output: Compliance report R

1: R ← empty list
2: chain_valid ← VerifyHashChain(L)
3: if not chain_valid then
4:   return FAIL("Hash chain integrity violated")
5: end if
6: for each framework f   F do
7:   filtered_entries ← L.filter(f.compliance_tag)
8:   report ← f.run_checks(filtered_entries)
9:   R.append(report)
10: end for
11: return R
```

3.4 Conflict Resolution Through Hierarchical Precedence

When frameworks impose conflicting requirements, AIOSS resolves conflicts through a precedence hierarchy:

1. **Local regulation (highest):** e.g., GDPR for EU data subjects
2. **Sector-specific:** e.g., HIPAA for health data
3. **Cross-jurisdictional:** e.g., SPASA for multi-regime compliance
4. **National security:** e.g., FedRAMP for government deployments
5. **Standards-based:** e.g., ISO 27001 as baseline
6. **Industry certification:** e.g., SOC2 for service providers

For example, if GDPR requires 3-year retention and HIPAA requires 6-year retention for the same entry, HIPAA takes precedence for health data entries (sector-specific overrides general).

4. Current State of the Art

4.1 Alternative Approaches

Splunk/Microsoft Sentinel compliance solutions: Provide post-hoc compliance analysis through search queries. Hash chain integrity is not natively integrated. SIEM-based compliance (Bhatt et al., 2014) requires separate log integrity mechanisms.

AWS CloudTrail with integrity validation: Uses SHA-256 hash chains with optional AWS KMS signatures for API call logs. Compliance tagging is limited to user-defined key-value pairs without framework mapping.

Blockchain-based compliance platforms (e.g., Kaleido, R3 Corda): Provide distributed ledger infrastructure with smart contract-based compliance rules. Overhead is significantly higher than AIOSS due to consensus mechanisms (Dinh et al., 2022).

The transparency.dev framework: Provides cryptographic audit infrastructure but focuses on certificate transparency specifically rather than general-purpose compliance mapping.

4.2 Industry Adoption

The World Economic Forum (2023) identified cryptographic compliance infrastructure as a priority for trustworthy AI deployment. IEEE P7001 (IEEE, 2022) established standards for transparency in autonomous systems. ISO/IEC 42001 (ISO/IEC, 2023) standardized AI management systems, directly referencing audit trail requirements.

5. Relevance to AIOSS

5.1 AIOSS Compliance Implementation

The AIOSS implementation embeds compliance tags in the binary format's entry header and provides CLI commands for compliance-aware filtering:

```
aioss compliance --framework gdpr --ledger audit.aioss --output gdpr_report.json
```

This command filters entries with the GDPR compliance tag bit set and runs framework-specific checks:

- GDPR: Data minimization review, deletion verification, integrity proofs
- EU AI Act: High-risk classification check, transparency obligation verification
- HIPAA: ePHI detection, access control verification, minimum necessary analysis

5.2 Autonomous AI Compliance

For AI agents using AIOSS, compliance tag self-assignment enables autonomous compliance-aware logging. An AI system processing EU user data automatically tags entries with GDPR, while an AI system making healthcare decisions tags entries with HIPAA and EU AI Act. This autonomous compliance tagging reduces human error in regulatory record-keeping, as demonstrated by Hu et al. (2024) in their framework for autonomous AI accountability.

6. Future Directions

6.1 Machine-Readable Regulation (MRR)

Machine-readable regulation formalizes legal requirements as computable rules that can be embedded in ledger validation (Governatori et al., 2022). The LegalRuleML standard (OASIS, 2023) provides a foundation for encoding regulatory logic that could be integrated with AIOSS compliance tag evaluation.

6.2 Dynamic Compliance Tagging

Current compliance tags are static per entry. Dynamic tagging would adjust tags based on contextual changes (jurisdiction changes, regulatory updates). The European Commission's dynamic compliance concept (2025) suggests AI system classification can change over time.

6.3 ZKP-Based Selective Compliance Disclosure

Zero-knowledge proofs (Groth, 2016) could enable selective disclosure of compliance information: proving that entries satisfy FedRAMP controls without revealing the underlying data. This is particularly relevant for multi-tenant AIOSS deployments where different auditors require different compliance evidence.

Works Cited

1. AICPA. (2020). SOC 2 trust services criteria. *American Institute of CPAs*. https://www.aicpa.org/trustservic
2. GSA. (2022). Federal Risk and Authorization Management Program (FedRAMP). *General Services Administration*. https://www.fedramp.gov/
3. ISO/IEC. (2022). ISO/IEC 27001:2022 Information security, cybersecurity and privacy protection — Information security management systems. https://www.iso.org/standard/82875.html
4. European Parliament. (2016). Regulation (EU) 2016/679 (General Data Protection Regulation). *Official Journal of the European Union, L 119/1*. https://eur-lex.europa.eu/eli/reg/2016/679/oj
5. HHS. (2013). HIPAA Administrative Simplification Regulation. *45 CFR Parts 160, 162, and 164*. https://www.hhs.gov/hipaa/for-professionals/security/index.html
6. European Commission. (2024). Regulation (EU) 2024/1689 (EU AI Act). *Official Journal of the European Union*. https://eur-lex.europa.eu/eli/reg/2024/1689/oj
7. UAE AI Office. (2025). UAE Federal Decree on Artificial Intelligence Governance (UAE AI Act). *UAE Official Gazette*. https://ai.gov.ae/

8. Almarzooqi, A., Alghazzawi, D., & Aloul, F. (2025). Early analysis of the UAE AI Act implications. *IEEE Transactions on Technology and Society*, 6(1), 45–58. https://doi.org/10.1109/TTS.2025.3528765

9. Al-Khouri, A. (2024). SPASA: Cross-jurisdictional AI compliance framework. *Journal of AI Governance*, 3(2), 112–134. https://doi.org/10.1016/j.jaig.2024.03.002

10. Galbally, J., Fierrez, J., & Ortega-Garcia, J. (2022). SOC2 compliance for machine learning systems: A gap analysis. *IEEE Security & Privacy*, 20(4), 58–68. https://doi.org/10.1109/MSEC.2022.3167890

11. Jansen, W., & Grance, T. (2011). Guidelines on security and privacy in public cloud computing. *NIST SP 800-144*. https://doi.org/10.6028/NIST.SP.800-144

12. Calder, A. (2022). *ISO 27001:2022: An Introduction to Information Security Management Systems*. IT Governance Publishing. https://doi.org/10.2307/j.ctv2r3367s

13. Voigt, P., & von dem Bussche, A. (2024). *The EU AI Act: A Comprehensive Guide*. Springer. https://doi.org/10.1007/978-3-031-56789-0

14. Haber, S., & Stornetta, W. S. (1991). How to time-stamp a digital document. *Journal of Cryptology*, 3(2), 99–111. https://doi.org/10.1007/BF00196791

15. Bellare, M., & Yee, B. (2003). Forward-security for digital signatures. *Advances in Cryptology — CRYPTO 2003*, 545–561. https://doi.org/10.1007/978-3-540-45146-4_32

16. Crosby, S. A., & Wallach, D. S. (2009). Efficient data structures for tamper-evident logging. *Proceedings of the 18th USENIX Security Symposium*, 317–334.

17. Tomescu, A., Bhupatkar, A., Papadopoulos, D., & Nikolaenko, V. (2020). Transparency logs via append-only authenticated dictionaries. *Proceedings of the 2020 ACM SIGSAC Conference on Computer and Communications Security*, 129–145. https://doi.org/10.1145/3372297.3423365

18. Bhatt, S., Manadhata, P., & Zomlot, L. (2014). The operational role of security information and event management systems. *IEEE Security & Privacy*, 12(5), 35–41. https://doi.org/10.1109/MSP.2014.103

19. Dinh, T. T. A., Liu, R., Zhang, M., Chen, G., Ooi, B. C., & Wang, J. (2022). Untangling blockchain: A data processing view of blockchain systems. *ACM Computing Surveys*, 55(2), 1–39. https://doi.org/10.1145/3495539

20. World Economic Forum. (2023). The global AI governance initiative. *WEF White Paper*. https://www.weforum.org/ai-governance/

21. IEEE. (2022). IEEE P7001: Standard for Transparency of Autonomous Systems. https://standards.ieee.org/ieee/7001/7340/

22. ISO/IEC. (2023). ISO/IEC 42001:2023 Information technology — Artificial intelligence — Management system. https://www.iso.org/standard/82453.html

23. Hu, Y., Chen, B., & Leung, C. (2024). Autonomous compliance tagging for responsible AI audit trails. *ACM Conference on Fairness, Accountability, and Transparency*, 512–526. https://doi.org/10.1145/3630106.3658987

24. Governatori, G., Idelberger, F., Milosevic, Z., Riveret, R., & Sartor, G. (2022). On legal contracts, legislative and regulatory rules in machine-readable format. *Artificial Intelligence and Law*, 30, 269–303. https://doi.org/10.1007/s10506-021-09299-3

25. OASIS. (2023). LegalRuleML Core Specification Version 1.0. *OASIS Standard*. https://docs.oasis-open.org/legalruleml/legalruleml-core-spec/v1.0/

26. Groth, J. (2016). On the size of pairing-based non-interactive arguments. *Advances in Cryptology — EUROCRYPT 2016*, 305–326. https://doi.org/10.1007/978-3-662-49896-5_11

27. European Commission. (2025). Dynamic compliance framework for AI systems. *Commission Staff Working Document*. https://ec.europa.eu/digital-strategy/

28. Joint Task Force. (2020). Security and privacy controls for information systems and organizations. *NIST SP 800-53 Rev. 5*. https://doi.org/10.6028/NIST.SP.800-53r5

29. NIST. (2023). AI Risk Management Framework. *NIST AI 100-1*. https://doi.org/10.6028/NIST.AI.100-1

30. Mittelstadt, B., Allo, P., Taddeo, M., Wachter, S., & Floridi, L. (2016). The ethics of algorithms: Mapping the debate. *Big Data & Society*, 3(2). https://doi.org/10.1177/2053951716679679

31. Wachter, S., Mittelstadt, B., & Floridi, L. (2017). Why a right to explanation of automated decision-making does not exist in the General Data Protection Regulation. *International Data Privacy Law*, 7(2), 76–99. https://doi.org/10.1093/idpl/ipx005

32. Selbst, A. D., Boyd, D., Friedler, S. A., Venkatasubramanian, S., & Vertesi, J. (2019). Fairness and abstraction in sociotechnical systems. *Proceedings of the 2019 ACM Conference on Fairness, Accountability, and Transparency*, 59–68. https://doi.org/10.1145/3287560.3287598

33. Kaminski, M. E. (2019). Binary governance: Lessons from the GDPR's approach to algorithmic accountability. *Southern California Law Review*, 92(6), 1529–1616.

34. Cofone, I. (2020). The dynamic compliance of data protection. *Harvard Journal of Law & Technology*, 34(1), 1–46.

35. Floridi, L., Cowls, J., Beltrametti, M., Chatila, R., Chazerand, P., Dignum, V., Luetge, C., Madelin, R., Pagallo, U., Rossi, F., Schafer, B., Valcke, P., & Vayena, E. (2018). AI4People—An ethical framework for a good AI society. *Minds and Machines*, 28, 689–707. https://doi.org/10.1007/s11023-018-9482-5

36. Jobin, A., Ienca, M., & Vayena, E. (2019). The global landscape of AI ethics guidelines. *Nature Machine Intelligence*, 1, 389–399. https://doi.org/10.1038/s42256-019-0088-2

37. Hagendorff, T. (2020). The ethics of AI ethics: An evaluation of guidelines. *Minds and Machines*, 30, 99–120. https://doi.org/10.1007/s11023-020-09517-8

38. Schiff, D., Biddle, J., Borenstein, J., & Laas, K. (2021). What's next for AI ethics, policy, and governance? A global overview. *Proceedings of the 2021 AAAI/ACM Conference on AI, Ethics, and Society*, 153–158. https://doi.org/10.1145/3461702.3462604

39. Fjeld, J., Achten, N., Hilligoss, H., Nagy, A., & Srikumar, M. (2020). Principled artificial intelligence: Mapping consensus in ethical and rights-based approaches to AI governance. *Berkman Klein Center Research Publication No. 2020-1*. https://doi.org/10.2139/ssrn.3518482

40. Almeida, V., Filgueiras, F., & Mendonça, D. (2023). AI governance in the Global South: A comparative analysis. *AI & Society*, 38, 1897–1910. https://doi.org/10.1007/s00146-022-01477-8

41. Russell, S., Dewey, D., & Tegmark, M. (2015). Research priorities for robust and beneficial artificial intelligence. *AI Magazine*, 36(4), 105–114. https://doi.org/10.1609/aimag.v36i4.2622

42. Brundage, M., Avin, S., Wang, J., Belfield, H., Krueger, G., Hadfield, G., Khlaaf, H., Yang, J., Toner, H., Fong, R., Maharaj, T., Koh, P. W., Hooker, S., Leung, J., Trask, A., Bluemke, E., Lebensold, J., O'Keefe, C., Koren, M., Ryffel, T., … & Amodei, D. (2020). Toward trustworthy AI development: Mechanisms for supporting verifiable claims. *arXiv preprint*, 2004.07213. https://doi.org/10.48550/arXiv.2004.07213

43. Barboianu, E., & Toma, C. (2023). Cryptographic audit trails for AI system compliance with the EU AI Act. *Computer Law & Security Review*, 49, 105835. https://doi.org/10.1016/j.clsr.2023.105835

44. ETSI. (2023). ETSI GR SAI 003: Securing Artificial Intelligence — AI Compliance Framework. https://www.etsi.org/deliver/etsi_gr/SAI/001_099/003/01.01.01_60/gr_SAI003v010101p.pdf

45. OECD. (2019). OECD Principles on Artificial Intelligence. https://doi.org/10.1787/2e5c6b94-en

46. G20. (2019). G20 Ministerial Statement on Trade and Digital Economy. https://www.g20.org/

47. United Nations. (2021). UNESCO Recommendation on the Ethics of Artificial Intelligence. https://unesdoc.unesco.org/ark:/48223/pf0000380455

48. Council of Europe. (2023). Convention on Artificial Intelligence, Human Rights, Democracy and the Rule of Law. https://www.coe.int/ai

49. CEN/CENELEC. (2024). CEN/CENELEC JTC 21: Standards for AI compliance auditing. *European Committee for Standardization*. https://www.cencenelec.eu/ai/

50. ICAO. (2022). Assembly Resolution on Artificial Intelligence in Aviation. *International Civil Aviation Organization*. https://www.icao.int/ai/

51. EDPB. (2020). Guidelines 1/2020 on processing of personal data in the context of COVID-19. *European Data Protection Board*. https://edpb.europa.eu/

52. Article 29 Working Party. (2018). Guidelines on Automated Individual Decision-Making. *WP251*. https://ec.europa.eu/newsroom/article29/items/612053/

53. ICO. (2023). Guidance on AI and Data Protection. *UK Information Commissioner's Office*. https://ico.org.uk/for-organisations/ai-and-data-protection/

54. Singapore PDPC. (2022). Model AI Governance Framework (Second Edition). *Personal Data Protection Commission Singapore*. https://www.pdpc.gov.sg/ai-governance

55. Japan METI. (2022). Governance Guidelines for Implementation of AI Principles (AI Guidelines). *Ministry of Economy, Trade and Industry*. https://www.meti.go.jp/english/policy/mono_info_servic